

Enforcing Truthful-Rating Equilibria in Electronic Marketplaces*

Thanasis G. Papaioannou, George D. Stamoulis
Department of Informatics, Athens University of Economics and Business
76 Patision Str., Athens, GR 10434, Greece
{pathan, gstamoul}@aueb.gr

Abstract

Reputation-based mechanisms and policies are vulnerable to the submission of untruthful ratings. In this paper, we define and analyze a game-theoretic model that captures the dynamics and the rational incentives in a competitive e-marketplace in which providers and clients exchange roles. We also study how we can enforce equilibria where ratings are submitted truthfully. We employ a mechanism prescribing that each service provision is rated by both the provider and the client, while this rating is included in the calculation of reputation only in case of agreement. First, we analyze the case where fixed monetary penalties are induced to both raters in case of disagreement. We prove that, under certain assumptions on the initial conditions, the system is led to a stable equilibrium where all participants report truthfully their ratings. We also investigate the introduction of non-fixed penalties to provide the right incentives for truthful reporting. We derive lower bounds on such penalties that depend on the participant's reputation values. Thus, by employing a punishment that is tailored properly for each participant, this approach can limit the unavoidable social welfare losses due to the penalties for disagreement.

1. Introduction

Electronic marketplaces such as eBay [5] have already become very popular for trading nearly all kinds of products and services. In such e-marketplaces there is hidden information on the quality and the type of the provider. Due to this asymmetry in information among transacting participants, there is a possibly high risk for a client to get less value from a transaction than expected. It is widely accepted that reputation can be a proper means of revealing the hidden information on low-performing participants in e-marketplaces as well as in peer-to-peer electronic environ-

ments that do not involve payments [3]. For such systems, we show in [10] that the calculation of reputation values has to be complemented by proper reputation-based policies that determine the pairs of peers eligible to interact with each other. However, the calculation of reputation is almost always based on the ratings of service provisions as submitted by the clients. Thus, reputation-based mechanisms are vulnerable to false or strategic feedback especially, if such reporting behavior provides some value to the rater. In [9], we introduce and analyze a mechanism that provides the incentives for truthful reporting in a peer-to-peer system. According to this mechanism, both transacting peers (rather than just the client) submit ratings on performance of their mutual transaction. If these are in disagreement, then both transacting peers are punished. The motivation for this is that such a disagreement is a sign that one of the transacting peers is lying, but the system cannot tell *who*, due to hidden information. Furthermore, the severity of each peer's punishment is determined by an extra metric, which is referred to as the peer's non-credibility metric; this metric is maintained by the mechanism and evolves according to the peer's record. On the other hand, if two transacting peers agree in their ratings, then this feedback is indeed taken into account in the calculation of the provider's reputation. As shown experimentally in [9], the combination of this mechanism with reputation-based policies is very effective even in peer-to-peer systems with dynamically evolving populations large fractions of which are lying.

In this paper, we deal with the problem of promoting truthful reporting of ratings in competitive e-marketplaces involving payments. This is a different context from peer-to-peer systems, which were considered in [9], and the problem is now more complicated due to the payments. That is, it is plausible that either the client or the provider has a direct monetary benefit from submitting a false rating and/or a future competitive advantage in the market, especially if providers and clients are able to exchange roles. We adapt properly the mechanism of [9]. In particular, in order to balance the risks due to hidden information of both the client and the provider of a transaction (i.e. a service- or a

*This work was partly supported by the Network of Excellence EuroNNGI (IST-2003-507613) funded by European Union.

product- provision), the former pays ahead a certain portion of the total payment, while the remaining amount is paid if both parties agree that this transaction was successful. A disagreement induces a direct monetary penalty to both parties involved, while gives rise to interesting trade-offs that also relate to the future positioning of the two parties in the market. The contribution of this paper is as follows: We develop and analyze a game-theoretic model capturing the above effects. We establish that with fixed penalty for disagreement our mechanism enforces truthful reporting under certain assumptions. We also investigate the case of non-fixed penalties depending on each participant's reputation value. We derive bounds for such penalties that do provide the right incentives for truthful reporting, while resulting in a reduction the social welfare losses that are due to unfair punishments. In the course of this analysis we also introduce an innovative reputation metric.

2. Related Work

Apart from [9], that has already been discussed in the introduction, there is significant published research that is related to truthful reporting. In [4], providers publish their services and their promised quality in a virtual market list. The proposed mechanism of [4] publishes a quality value for next service so as to compensate for the payoff gains resulted by deviations between the promised and the offered quality. In case of truthful clients, truthful reporting by providers is a subgame perfect equilibrium.

A side payment approach for eliciting honest feedback in electronic markets is proposed by Miller et al. in [6]. In particular, a payment charged to the current buyer is paid to the subsequent buyer according to a scoring rule for her prediction of the rating of a later buyer for their common seller. This approach enables truthful reporting of clients to be a Nash equilibrium, yet not a unique one. Another side payment mechanism is proposed by Jurca and Faltings in [7] for truthfully reporting on the performance of a hotel in the online hotel booking industry. The occupancy of the hotel is a linear function of its reputation. After a service provision, the hotel reports whether having offered high quality or not. Then, the client decides on agreeing or disagreeing with the report of the hotel, or leaving. This situation corresponds to a *Stackelberg* game. In case of disagreement, different fixed punishment fees ϵ_H and ϵ_C , with $\epsilon_H > \epsilon_C$ are charged to the hotel and the client respectively, and a negative vote for the hotel is counted. Otherwise, a positive vote is counted for the hotel. It is proved that when the hotel and the clients are rational, the percentage of false reports is bounded. Also, it is proved that when there is employed reputation for truthful reporting and clients are long-lived, then the number of cheats against the same client are bounded. Strategic reporting by clients was not considered

in [7], which is a major consideration in our work both in [9] and in this paper. Finally, a public goods' model is considered by Hauert et al. in [2]: population members may contribute or not to the public good and they may exert or not some effort to punish defectors. It is proved in [2] that the "a-social" equilibrium (i.e. all players do not contribute and do not punish) is the only evolutionary stable strategy (see Section 4). However, if reputation for being a punisher is introduced then the "social" equilibrium (i.e. all contribute and punish) has a significant basin of attraction.

In this paper, we cover several issues on truthful reporting of ratings' feedback that either are innovative or are not covered adequately by the aforementioned works. Such issues are: a) The stability of the truthful-rating Nash equilibrium that is enforced by the mechanism. b) The context of our study, where participants can act both as providers and as clients. c) The competition among participants on the basis of reputation (due to reputation-based policies) and its impact on the incentives regarding honesty in feedback reporting. d) The introduction and the study of the effectiveness of tailored reputation-based punishment for untruthful reporting, aiming to limit the social welfare losses induced by the penalties for disagreement.

3. The Basic Model

We consider an e-marketplace with many participants, where each of them can act both as a provider or as a client of a certain service (or product). There are many interesting service examples that fit this case. Consider, e.g., an e-marketplace for exchanging (in the sense of selling and buying) software modules among programmers, computational resources such as CPU and storage capacity, or even virtual communities of users that exchange some collectable good (e.g. vinyl records of classic music, stamps etc.). Next, we present our assumptions to model such an e-marketplace. Time is discretized in rounds. At the beginning of each round, each of the N participants decides whether she will act as a provider (with probability q) or as a client (with probability $1-q$) to select a provider, where q is bounded away both from 0 and 1. N is taken to be large and fixed. The model to be presented below, as well as the analysis, apply to both: a) a static environment with a fixed set of participants, and b) a dynamic one, where each participant has a geometrically distributed lifetime with expected value T rounds, at the end of which she is replaced by new one. An instance of the service (or the product) exchanged generates a fixed utility u to the client regardless of who was the provider, yet under the condition that the service was provided successfully, i.e. with quality above a pre-specified level. Each successful service provision demands costly effort v to be exerted by the provider. The participants of the market may belong to different performance

types. In particular, associated with each participant i is a success probability a_i , which is private information known only to himself. A reputation metric for performance is employed in order to reveal performance types of the participants. After each service provision, *both* participants are expected to submit feedback to the reputation system on the outcome of the service. This feedback includes a binary characterization of the service (i.e. “successful” or “unsuccessful”), combined with a measurable parameter whenever possible, in order to justify the binary rating; e.g. in the case of trading storage capacity the participants that transacted may also report the Tbytes provided. (This mechanism requires certain observing authorities and can be either centralized or distributed; see [9].) If the two feedback reports are in agreement, then the corresponding binary rating is aggregated into the reputation value of the provider. This is attained by employing an appropriate aggregation function such as Beta [1] or a new one that we introduce later; see Subsection 5.1. If either the feedback reports are not consistent or exactly one of the parties does not report feedback, then a monetary penalty c is induced to both parties. Note that inclusion of the measurable parameter in the feedback report prevents the participants that transacted from colluding *temporarily* in order to agree strategically. Furthermore, each service instance is charged at a fixed price b in all rounds. This price is set by the market and does not depend on the specific participants that transact. Thus, due to hidden information, the following risks arise in a service provision: a) the client may receive a service of unsatisfactory quality, while b) the provider may provide a service successfully and not be paid by the client. In order to balance these risks the client *prepays* a certain portion $p \cdot b$ of the total price b , while the remaining amount $(1 - p) \cdot b$ is paid only if both parties agree that this transaction was successful. Refusal of post-payment by the client is accompanied with a negative rating.

In order for the right incentives for performance to be provided to participants we assume that a reputation-based policy for provider selection is employed [10]. In particular, in each round, clients are assumed to associate to providers in a probabilistically fair manner according to the reputation values of the latter [10]. That is, if participant i serves at a certain round as a provider, then the probability that she will be selected by a certain participant j who acts now as a client is proportional to r_i . Of course, multiple or even all clients can select and be served by the same provider, contrary to what was assumed for peer-to-peer systems in [10]. The other provider selection policy analyzed in [10] is Max-Reputation, which prescribes that each client simply selects the provider with the highest reputation. If, however, we assume that only a uniformly picked random subset of the providers can offer what a certain client demands in each round, then Max-Reputation should be restricted to the ap-

propriate subset of providers. Under this assumption, the demand attracted by a participant when acting as a provider tends to approximate that of the assumed probabilistically fair selection reputation-based policy. Furthermore, say that in a certain round, the set of participants to act as clients (resp. providers) is C (resp. P). Note that the expected size of the set C of clients is $E[|C|] = N(1 - q)$, while that of P is $E[|P|] = Nq$. Since N is large and q is bounded away from 0 and 1, both sets of providers and clients are expected to have a large number of members selected randomly. Therefore, we have $\frac{1}{|P|} \sum_{k \in P} r_k \approx \frac{1}{N} \sum_{k=1}^N r_k = \bar{r}$, which implies that $\sum_{k \in P} r_k \approx \bar{r}|P| \approx \bar{r}Nq$, where \bar{r} is the mean reputation value over all participants in the market. Furthermore, notice that the probability for a certain provider i to be selected by a certain client j is $\frac{r_i}{\sum_{k \in P} r_k}$. Using the previous approximation, it follows that this probability equals approximately to $R_i/(Nq)$, where $R_i = \frac{r_i}{\bar{r}}$ is referred to as *rank* of participant i . In other words, the demand attracted by provider i depends on her rank R_i rather than directly on her reputation value r_i .

In Figure 1, we present the payoff matrix of the single-shot game arising, which is played in each round by every participant acting as a client with the provider she has selected. This game consists of two sub-games, namely that for the case of success of the service provided and that for the case of failure. We assume that the parties involved have three alternative pure reporting strategies, namely those in $S = \{Witness, Lie, Duck\}$. Thus, they can choose either to submit feedback truthfully or to submit dishonest feedback or not to send feedback at all. If the players’ feedback ratings are in disagreement, or if exactly one of them does not submit feedback, then no rating is taken into account but both players incur a monetary penalty c . Note that, to keep the present basic model simpler yet indicative of the main incentives’ issues arising, we have taken that the provider and the client pay the same penalty. This restriction is relaxed in the extended model of Section 5. On the other hand, in case of agreement, the common rating is taken into account and the provider’s reputation is updated. Thus, due to the reputation-based policy described above, each agreed rating affects the future payoffs of the players. In particular, the impact of a positive rating in the future payoff for a provider is denoted as $w_p > 0$ and that of a negative rating is denoted as $-w'_p < 0$. The payoff of a client is affected too by an agreed rating, since the rating influences the mean reputation \bar{r} , and in turn the client’s rank, which determines the demand that she will attract whenever she acts as a provider in the future. The payoff impact for a client of an agreed positive rating is denoted as $-w_c < 0$ and that of a negative one is denoted as $w'_c > 0$. The payoff impacts $w_p, -w'_p, w_c, -w'_c$ are taken as fixed, i.e. independent of the provider’s and the client’s rank values, which as already explained influence the respective demand they will

		client			
		Witness	Lie	Duck	
provider	success a_i	Witness	$u-b-w_c$	$u-p-b-c$	$u-p-b-c$
		Lie	$u-b-c$	$u-p-b+w_c'$	$u-p-b-c$
		Duck	$u-b-c$	$u-p-b-c$	$u-p-b$
	fail $1-a_i$	Witness	$-p-b+w_c'$	$-p-b-c$	$-p-b-c$
		Lie	$-p-b-c$	$-p-b-w_c$	$-p-b-c$
		Duck	$-p-b-c$	$-p-b-c$	$-p-b$

Figure 1. The single-shot game.

attract in the future. This assumption is relaxed in Section 5. It can be seen that a system where the above payoff impacts are indeed almost fixed is a large market employing a reputation measure taking infinitely increasing/decreasing values; e.g. when reputation equals the difference between successful and unsuccessful service provisions of a certain provider. Thus, the game defined captures the trade-offs for each of the players to submit a truthful feedback, versus to lie or to not submit feedback at all. As already explained, all terms of the payoff matrix of the game are *independent* of the players' reputation values. Moreover, each of the players decides on her strategy after observing the outcome of the service provision, which is much more revealing information than the reputation of the provider. The *control parameter* of the mechanism that can influence the players' reporting behavior is the penalty c in case of disagreement, and to a lesser extent the portion p of the price that is pre-paid. All other terms of the payoff matrix are determined by the way the market operates and evolves.

4. Stability Conditions for Truthful Equilibrium

In this section, our goal is to find necessary conditions under which truthful reporting by participants in both roles is a stable Nash equilibrium in both the success and the failure subgames of Figure 1. Therefore, we first employ equilibrium analysis for the two subgames of Figure 1, and find conditions so as truthful reporting to be best response to itself both for providers and for clients regardless of the outcome of the transactions. Then, we employ evolutionary game theory [8], in order to study the stability of the desired Nash equilibrium. To the former end, and parallel-

ing part of the analysis of [7], if participants in the client role could be assumed to be truthful, it is a dominant strategy for providers to be truthful as well when they succeed in providing their services. In case that they fail, truthful reporting can be again dominant strategy for providers, if the monetary punishment that is charged to providers and clients in case of disagreement is $c > w_p'$. Thus, when clients are truthful, providers can be made truthful too. On the other hand, if providers could be assumed to be truthful, it is dominant strategy for clients to be truthful as well when providers fail in service provision. When providers succeed, it should be that $c > b(1-p) + w_c$, in order for truthful reporting to be dominant strategy for clients too. Summarizing the above discussion, it should be that $c > b(1-p) + w_c$ and $c > w_p'$ in order for *Witness* to be best response to itself. These conditions are henceforth assumed.

The only reasonable assumption regarding reporting is that both providers and clients act rationally so as to maximize their payoffs. Henceforth, we analyze our game-theoretic model under this assumption. First, note that there is no disagreement punishment c (and payoff impacts of votes w_p, w_c, w_p', w_c') that render *Witness* dominant strategy in either subgame of Figure 1. When $c > b(1-p) + w_c$ and $c > w_p'$, there are three pure Nash equilibria and one in mixed strategies in each subgame: i) Both provider and client report truthfully. ii) Both provider and client lie. iii) Both provider and client do not submit feedback. iv) Mixed strategies for transacting parties for either subgame, denoted as x_M, \tilde{x}_M for providers in the success and the failure subgame respectively, and as y_M, \tilde{y}_M for clients in the success and the failure subgame respectively, and depicted in Figure 2. Two important questions are in order: Which Nash equilibrium is ultimately selected by participants? Is this equilibrium stable? In order to answer these, we employ the theory for evolutionary stability in games [8]. According to this theory, an *Evolutionary Stable Strategy* (ESS) is a Nash equilibrium strategy that is also a better response to any mutant strategy played by a small fraction of invaders than the mutant strategy is to itself. A totally or partly dominated strategy cannot be ESS. However, since $c > b(1-p) + w_c$ and $c > w_p'$, no dominated strategies exist. According to Proposition 2.4 of Samuelson [8], in an asymmetric normal-form game (such as the one of Figure 1), $\sigma^* = (\sigma_I^*, \sigma_{II}^*)$ is an evolutionary stable strategy (ESS) of its *symmetric* version (i.e. one that designates actions for a participant in both of her roles in the game) if and only if $(\sigma_I^*, \sigma_{II}^*)$ is a strict Nash equilibrium of the asymmetric game, with $\sigma_I, \sigma_{II} \in \Delta S$. Thus, all four aforementioned equilibria of each subgame can be ESS.

Next, we investigate the necessary conditions in order for the [*Witness, Witness*] strategy-pair to be selected by participants in either subgame, i.e. to be the ESS that the e-marketplace will evolve to. For this purpose, we study

the relative fitnesses (i.e. expected payoffs) of strategies and their evolutionary dynamics [8], i.e. the marginal difference along time of the fractions of the participants that play the three strategies in either subgame. According to [8], the evolutionary dynamics for a pure strategy s with expected payoff π_s that is played by a population fraction x_s (or alternatively when all play a mixed strategy that assigns probability x_s to strategy s) are given by:

$$\dot{x}_s = x_s(\pi_s - \bar{\pi}), \quad (1)$$

where $\bar{\pi} = \sum_{s'} x_{s'} \pi_{s'}$ is the average expected payoff over the various pure strategies the game. Thus, the evolutionary dynamics in the game of Figure 1 are given by: $\dot{x}_1 = x_1(\pi_{s_1} - \bar{\pi})$, $\dot{x}_2 = x_2(\pi_{s_2} - \bar{\pi})$, $\dot{x}_3 = x_3(\pi_{s_3} - \bar{\pi})$, $\dot{y}_1 = y_1(\hat{\pi}_{s_1} - \hat{\pi})$, $\dot{y}_2 = y_2(\hat{\pi}_{s_2} - \hat{\pi})$, $\dot{y}_3 = y_3(\hat{\pi}_{s_3} - \hat{\pi})$. where s_1, s_2, s_3 stand for *Witness*, *Lie*, and *Duck* respectively. These definitions apply to both subgames. For example, for a provider that plays *Witness* in the success subgame $\pi_1 = y_1(b - v + w_p) + (y_2 + y_3)(pb - v - c)$; also $\bar{\pi} = x_1 \cdot \pi_1 + x_2 \cdot \pi_2 + x_3 \cdot \pi_3$. Finally, the conditions under which a strategy is ESS amounts to deriving its corresponding *basin of attraction*. The basin of attraction for the [*Witness*, *Witness*] strategy-pair corresponds to a region for (x_1, x_2, x_3) and (y_1, y_2, y_3) with the following property: If the system's population mix lies in this region at some point in time, then it ultimately evolves towards $(1, 0, 0)$ and $(1, 0, 0)$, which implies that all participants will asymptotically (in time) be reporting truthfully both as providers and as clients in either subgame. Of course, a different basin of attraction is associated with each of the other Nash equilibria. Since $x_1 = 1 - x_2 - x_3$ and $y_1 = 1 - y_2 - y_3$ each basin of attraction can be expressed as the Cartesian product of two two-dimensional regions.

Proposition 1. *The basin of attraction of ESS truthful reporting is the region $X^* \times Y^*$, where $(x_2, x_3) \in X^*$ iff $x_2 < \min\{\frac{c-b(1-p)-w_c-(c-w_c)x_3}{2c+w'_c}, 1 - \frac{2c+w'_c}{c+w'_c}x_3\}$ and $x_3 < \frac{c+w'_c}{2c+w'_c}$, and $(y_2, y_3) \in Y^*$ iff $y_2 < \min\{\frac{c-w'_p}{2c-w'_p+w_p}(1-y_3), 1 - \frac{2c-w'_p}{c-w'_p}y_3, 1 - \frac{2c+w_p}{c+w_p}y_3\}$ and $y_3 < \min\{\frac{c-w'_p}{2c-w'_p}, \frac{c+w_p}{2c+w_p}\}$.*

Outline of Proof. Being a strict Nash equilibrium of the game, the strategy-pair [*Witness*, *Witness*] is also an ESS if and only if the truthful reporting strategy is more fit than the other two pure reporting strategies of providers and clients. Due to Equation (1), truthful reporting being more fit is equivalent to the following conditions in both subgames:

$$\frac{\dot{x}_1}{x_1} > \frac{\dot{x}_2}{x_2}, \frac{\dot{x}_1}{x_1} > \frac{\dot{x}_3}{x_3}, \frac{\dot{y}_1}{y_1} > \frac{\dot{y}_2}{y_2}, \frac{\dot{y}_1}{y_1} > \frac{\dot{y}_3}{y_3}, \quad (2)$$

or equivalently to: $\pi_{s_1} > \pi_{s_2}$, $\pi_{s_1} > \pi_{s_3}$, $\hat{\pi}_{s_1} > \hat{\pi}_{s_2}$, $\hat{\pi}_{s_1} > \hat{\pi}_{s_3}$ in both subgames. Using the expressions for the various payoffs, we obtain the conditions defining X^* and Y^* after some algebra. \square

Figure 2 depicts the basin of attraction for truthful reporting for the provider role (left part) and for client role (right part), taking that $w_p > w'_p$. See Section 6 for numerical values of the various parameters. When $(x_2, x_3) \times (y_2, y_3) \in X^* \times Y^*$ at some point in time, then the evolutionary dynamics ultimately lead all participants to truthful reporting in both cases of success or failure in service provision.

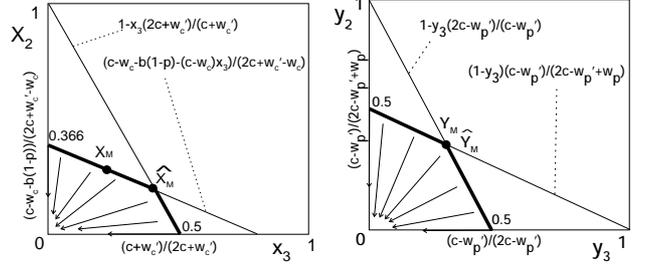


Figure 2. The basin of attraction sought.

5. Extended Model: Reputation-based Payoff Impact and Punishments

In this section, we extend the model of Section 3 by calculating the payoff impact resulted by a positive or negative vote for participants on the basis of reputation (essentially of the rank), and by analysing the resulting repeated game.

5.1. Computing Reputation

Feedback should be aggregated properly in order the history of a provider's performance to be summarized in a meaningful reputation metric. Beta aggregation function [1] has been proved to accurately update the posterior probability that a provider offer its services successfully. Moreover, if proper reputation-based policies [10] are employed, then the resulting reputation provides incentives to providers for increased performance. According to Beta aggregation function, the posterior belief on the success probability of a provider is the fraction of the number of its successful service provisions z over its total service provisions n with more weight being placed to the recent history. Thus, following a new service provision (i.e. $n \rightarrow n + 1$), reputation of the provider is updated as follows: $r' = \frac{\beta z + 1(\text{success})}{\beta n + 1}$, with β being the discount factor. However, this function for reputation estimation results to *time-dependent* impact of a single vote to the reputation and rank values of the transacted provider and client, which complicates our analysis. For example, the difference in rank for a provider with previous reputation $r = \frac{z}{n}$ and rank R after a positive vote is $\Delta R_p^+ = \frac{r'}{\bar{r} - \frac{r}{N} + \frac{1}{N}} - R = \frac{(\beta r n + 1)N}{\bar{r}N(\beta + 1) + 1 - r} - R$, where we

have taken into account how r and \bar{r} are updated. Notice that ΔR_p^+ depends on the number of transactions n , which in turn depends on the time of the provider in the system. Recall now that the demand to be attracted by a participant in the future depends on her rank. Therefore, we need a reputation metric for which the resulting updating of the rank should be time-independent.

For this purpose, we introduce a *new* more tractable reputation metric that is updated according to the formula $r' = \beta r + (1 - \beta)\mathbf{1}(\text{success})$. This aggregation function provides reputation values equal to the numerator of Beta. Thus, in order to maintain a high reputation value (i.e. close to 1), a provider should keep on offering services successfully. We have evaluated this reputation metric when applied in conjunction with the reputation-based policies of [10] by means of simulation experiments similar the ones presented in [10]. It turns out that, similarly to the reputation metric obtained by means of Beta, this new metric provides incentives to participants for high performance in service provision. These experiments are omitted for brevity reasons. Thus, we henceforth employ the proposed reputation function in the subsequent calculations of the paper. The difference in the ranks of the transacted participants after a positive vote are denoted as ΔR_p^+ , ΔR_c^+ for the provider and the client respectively. The respective differences after a negative vote are ΔR_p^- , ΔR_c^- . In particular: $\Delta R_p^+ = \frac{\beta R \bar{r} + 1 - \beta}{\bar{r} + \frac{1 - \beta}{N}(1 - R \bar{r})} \cdot R$, $\Delta R_p^- = \frac{\beta R}{1 - R \frac{1 - \beta}{N}} \cdot R$, $\Delta R_c^+ = \frac{r}{\bar{r} + \frac{1 - \beta}{N}(1 - \bar{r})} \cdot R$, $\Delta R_c^- = R \frac{1 - \beta}{1 - \frac{1 - \beta}{N}} \cdot R$. All these differences are independent of the number of service provisions already performed by the provider, which is the desired property that simplifies the derivation of reputation-based punishments presented in the next subsection.

5.2. The Repeated Game

Henceforth, we consider the single-shot game of Figure 1 played *repeatedly*, yet with two important differences: a) In case of disagreement, the monetary punishment depends on the rank of the participant as well as on whether she acted as provider or client. b) In case of agreement, the future-payoff impacts $w_p, -w_p', w_c, -w_c'$ are not included in the total payment as fixed terms. Instead, we quantify them precisely as follows: We take explicitly into account the impact of the agreed vote on the ranks of each of the two participants in her future expected payoff for each of the rounds to follow. We assume that the payoff of each extra future round is discounted by an extra factor δ . In fact, in the case of dynamically evolving population, $\delta = d/T$ where d is the discount factor and T is the expected lifetime of each participant. We thus derive rank-dependent monetary punishments that render truthful reporting a Nash equilibrium of the repeated game. First, we compute the

per round expected payoff of a participant assuming that all participants report truthfully. In particular, at round t a participant i with rank $R_i^{(t)}$ and success probability a_i , has an expected per round payoff $V(R_i^{(t)}, a_i)$ that equals:

$$V(R_i^{(t)}, a_i) = q \frac{1 - q}{q} R_i^{(t)} [a_i(b - v) + (1 - a_i)pb] + (1 - q)[\bar{a}(u - b) + (1 - \bar{a})(-pb)]. \quad (3)$$

This expectation applies prior to the randomized decision of the participant regarding whether she will act as a provider or a client. In equation (3), \bar{a} is an estimate of the probability that the participant will be provided a successful service by the ‘‘average’’ provider if she acts as a client in the round considered. Next, we derive the reputation-based punishments in the case of disagreement.

Proposition 2. *Truthful reporting is a Nash equilibrium of the repeated game for all players in both roles of the game, if monetary punishments $c_i > f(R_i)$ and $\hat{c}_j > \hat{f}(R_j)$ are charged in case of disagreement in feedback reports to transacted parties i (provider) and j (client) with ranks R_i and R_j respectively; $f(\cdot)$, $\hat{f}(\cdot)$ are functions of rank.*

Outline of Proof. According to [8], truthful reporting is a Nash equilibrium of the repeated game if and only if even a deviation from this strategy for a single round assigns to the deviant participant a strictly smaller payoff than continuous truthful reporting. Clearly, such a deviation from truthful reporting is only meaningful for a provider that failed in service provision and for a client that was offered a service successfully. For brevity reasons, we only present the analysis of the case of a provider that failed. The case of a client is treated similarly. We consider a particular participant i that acted unsuccessfully as a provider at round t . Since the corresponding client is by assumption truthful, if i would either ‘‘Lie’’ or ‘‘Duck’’ at the end of round t , then she would incur a punishment c_i while her rank $\tilde{R}_i^{(t+1)}$ at the next round would equal $R_i^{(t)}$. On the other hand, if participant i reports truthfully, then her future rank $R_i^{(t+1)}$ will be $R_i^{(t+1)} = R_i^{(t)} + \Delta R_p^- < R_i^{(t)}$. Thus, by being truthful, participant i avoids the immediate loss due to the monetary punishment but experiences a loss in her future payoff, which equals $\sum_{\tau=t+1}^{\infty} \delta^{\tau-t} [V_i(\tilde{R}_i^{(\tau)}, a_i) - V(R_i^{(\tau)}, a_i)]$, where: $\{R_i^{(\tau)}\}_{\tau=t+1}^{\infty}$ is the sequence of future ranks of participant i if she is truthful at round t and all other rounds to follow while $\{\tilde{R}_i^{(\tau)}\}_{\tau=t+1, \dots}$ is the sequence of future ranks of participant i if she is not truthful at round t (thus avoiding the reduction of her future rank) and she is truthful in all other rounds to follow. Therefore, the deviation under discussion will not occur, as long as the punishment c_i is large enough

to exceed the loss in future payoffs, i.e. if and only if

$$c_i > \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} [V(\tilde{R}_i^{(\tau)}) - V(R_i^{(\tau)})]. \quad (4)$$

Using the expression of equation (3) for the expected per round payoff, it follows that $V(\tilde{R}_i^{(\tau)}, a_i) - V(R_i^{(\tau)}) = (1-q)E[\tilde{R}_i^{(\tau)} - R_i^{(\tau)}] \cdot [a_i(b-v) + (1-a_i)pb]$. Notice that since N is large we can approximate the rank differences for the provider and the client by omitting the terms $\Theta(\frac{1}{N})$; i.e. $\Delta R_p^+ \approx (1-\beta)(\frac{1}{\bar{r}} - R_i)$, $\Delta R_p^- \approx -(1-\beta)R_i$ and $\Delta R_c^+ = \Delta R_c^- \approx 0$. Therefore, we obtain after some algebra, $\tilde{R}_i^{(\tau)} - R_i^{(\tau)} \approx (1-q+q\beta)^{\tau-t-1}(-\Delta R_p^-) = (1-q+q\beta)^{\tau-t-1}(1-\beta)R_i^{(t)}$. To this end, the condition of equation (4) for the punishment c_i of the participant i that acted as provider becomes

$$\begin{aligned} c_i > \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} (1-q)(1-q+q\beta)^{\tau-t-1} (-\Delta R_p^-) \\ & [a_i(b-v) + (1-a_i)pb] \approx \\ & \approx (1-q) \frac{\delta(1-\beta)}{1-\delta(1-q+q\beta)} R_i^{(t)} [a_i(b-v) + (1-a_i)pb]. \end{aligned}$$

This completes the derivation of the rank-dependent punishment in case of disagreement for the participant that acted as a provider. \square

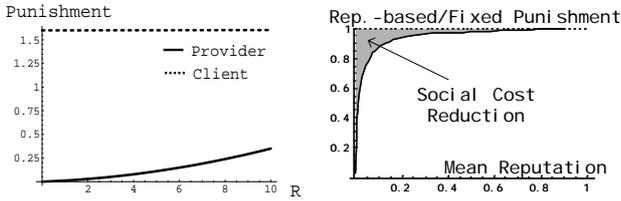


Figure 3. Left: Penalty for disagreement for participants w.r.t. their ranks when $\bar{r} = 0.1$. Right: Reduction of social cost due to punishment w.r.t. mean reputation.

6. Numerical Example - Future Work

Consider an electronic marketplace according to the model of Section 3. The various parameters of the model are: $N = 1000$, $q = 0.4$, $p = 20\%$, $b = 2$, $u = 2.5$, $v = 0.5$, $\beta = 0.6$. For a participant with rank $R = 1$ and $a \approx 0.5$ if $\bar{r} = 0.5$, according to Section 5, the values of the various payoff impacts are $w_p = \frac{\delta}{1-\delta} \Delta R_p^+ g(R, a, s_1) = 0.02251$, $-w'_p = \frac{\delta}{1-\delta} \Delta R_p^- g(R, a, s_1) = -0.02253$,

$-w_c = \frac{\delta}{1-\delta} \Delta R_c^+ \hat{g}(R, a, s_1) = -0.00002$ and $w'_c = \frac{\delta}{1-\delta} \Delta R_c^- \hat{g}(R, a, s_1) = 0.00002$, where $g(R, a, s_1) = q \frac{1-q}{q} R [a(b-v) + (1-a)pb]$. The basins of attraction for clients and providers for these parameters and for punishment for disagreement $c = 6$ are depicted in Figure 2. The reputation-based punishments with respect to R when $\bar{r} = 0.5$ for a provider and for a client are depicted in the left part of Figure 3. Assuming a Normal distribution of ranks in the electronic marketplace with $(\mu, \sigma) = (1, 0.5)$, reputation-based punishment results in a reduction of the social cost paid for punishments in comparison with fixed punishment for all mean reputation values (see the shaded region in the right part of Figure 3). As a future work, we plan to derive bounds in the achievable reduction of the social cost by reputation-based punishment. Also, we will analyse the evolutionary stability of the truthful equilibrium with reputation-based punishment.

References

- [1] A. Jøsang, S. Hird, E. Faccer. Simulating the effect of reputation systems on e-markets. In *Proc. of the 1st International Conference on Trust Management*, Crete, Greece, May 2003.
- [2] C. Hauert, N. Haiden, and K. Sigmund. The dynamics of public goods. *Discrete and Continuous Dynamical Systems - Series B*, 4(3):575–587, August 2004.
- [3] C. Dellarocas. *Reputation Mechanisms*. Forthcoming in Handbook on Economics and Information Systems (T. Hendershott, ed.), Elsevier Publishing. Available at: <http://faculty.haas.berkeley.edu/hender/ISEcon/ISEcon.htm>.
- [4] C. Dellarocas. Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality. In *Proc. of the Workshop on Agent-Mediated Electronic Commerce*, July 2002.
- [5] C. Dellarocas. Sanctioning reputation mechanisms in online trading environments with moral hazard. MIT Sloan Working Paper No. 4297-03, July 2004. Available at SSRN: <http://ssrn.com/abstract=393043>.
- [6] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Forthcoming in Management Science*, 2006. Available at: <http://www.si.umich.edu/presnick/papers/elicite/>.
- [7] R. Jurca, B. Faltings. An incentive compatible reputation mechanism for the online hotel booking industry. In *Proc. of IEEE Conference on E-Commerce*, Newport Beach, SA, USA, May 2004.
- [8] L. Samuelson. *Evolutionary Games and Equilibrium Selection*. The MIT Press, 1997. ISBN 0-262-19382-5.
- [9] T. G. Papaioannou, G. D. Stamoulis. An incentives' mechanism promoting truthful feedback in peer-to-peer systems. In *Proc. of IEEE/ACM CCGRID (Workshop on Global P2P Computing)*, Cardiff, UK, May 2005.
- [10] T. G. Papaioannou, G. D. Stamoulis. Reputation-based policies that provide the right incentives in peer-to-peer environments. *Computer Networks, Special Issue on Management in Peer-to-Peer Systems*, 50(4):563–578, March 2006.